The Patent Office
COPYRIGHT • TRADE MARKS

INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

CERTIFIED COPY OF PRIORITY DOCUMENT

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 26 July 2001

THIS PAGE BLANK (USPTO)

Patents Form 1/77

Patents Act 1977
(Rule 16)

27 JUL 2000

The Patent Office

27JUL00 E556041-1 D00393
P01/7700 0.00-0018328.5

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1.	Your reference	2000P04909/GB/R76/DA/cs
2.	Patent application number *(The Patent Office will fill in this part)*	27 JUL 2000 **0018328.5**
3.	Full name, address and postcode of the or of each applicant *(underline all surnames)*	Roke Manor Research Limited Roke Manor Research Old Salisbury Lane Romsey, Hampshire SO51 0ZN
	Patents ADP number *(if you know it)*	5615455005
	If the applicant is a corporate body, give the country/state of its incorporation	Great Britain
4.	Title of the invention	CENTRALISED BANDWIDTH ALLOCATION FOR MULTICAST TRAFFIC
5.	Name of your agent *(if you have one)*	DEREK ALLEN
	"Address for service" in the United Kingdom to which all correspondence should be sent *(including the postcode)*	Siemens Shared Services Limited Intellectual Property Department Siemens House, Oldbury Bracknell, Berkshire RG12 8FZ United Kingdom
	Patents ADP number *(if you know it)*	2898443006
6.	If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and *(if you know it)* the or each application number	Country / Priority application number *(if you know it)* / Date of filing *(day / month / year)*
7.	If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	Number of earlier application / Date of filing *(day / month / year)*
8.	Is a statement of inventorship and of right to grant of a patent required in support of this request? *(Answer 'Yes' if:* *a) any applicant named in part 3 is not an inventor, or* *b) there is an inventor who is not named as an applicant, or* *c) any named applicant is a corporate body.* *See note (d))*	YES

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form	
Description	4
Claim *(s)*	
Abstract	
Drawing *(s)*	

10. If you are also filing any of the following, state how many against each item.

Priority documents	
Translations of priority documents	
Statement of inventorship and right to grant of a patent *(Patents Form 7/77)*	
Request for preliminary examination and search *(Patents Form 9/77)*	
Request for substantive examination *(Patents Form 10/77)*	
Any other documents *(please specify)*	

11. I/We request the grant of a patent on the basis of this application.

Signature Date 25/07/2000

Derek Allen

12. Name and daytime telephone number of person to contact in the United Kingdom — Derek ALLEN - 01344 396808

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

a) *If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.*

b) *Write your answers in capital letters using black ink or you may type them.*

c) *If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.*

d) *If you have answered 'Yes' Patents Form 7/77 will need to be filed.*

e) *Once you have filled in the form you must remember to sign and date it.*

f) *For details of the fee and ways to pay please contact the Patent Office.*

Centralised Bandwidth Allocation for Multicast Traffic

Background

Traffic volume in the Internet is growing exponentially, doubling every 3-6 months. The current capacity of IP Routers is insufficient to meet this demand and hence there is a market opportunity for products that can route IP traffic at extremely large aggregate bandwidths in the order 10 Gbit/s - several Terra bit/s. Such routers are termed "Super Routers".

Additionally, there is a growing need to support multicast (one to many or many to many) communications within the internet or any other IP based network. To support such service an IP router must be able to replicate packets and send them to multiple outputs on a per packet basis. In a router where bandwidth allocations are strictly controlled (in order to support Quality of Service criteria) it is necessary to determine how much bandwidth to allocate to multicast traffic across the core switching fabric.

Invention

The invention relates to a method of dynamically adjusting the bandwidth, allocated to multicast traffic, across an ATM or crossbar like switching fabric that joins several IP packet forwarder functions to form a "Super Router" node.

Such a scheme is presented in the following description. In order to prevent head of line blocking unicast traffic is queued in separate logical scheduling entities (called scheduler blocks) according to which egress forwarder it is destined. The scheduler block serves a set of queues (per class or per connection) via any mechanism desired (e.g. strict priority or Weighted Fair Queuing) provided that the real time IP traffic class is guaranteed a minimum bandwidth.

However, for multicast traffic it is not practical to queue traffic on the basis of a unique combination of egress destinations as the number of queues required becomes unmanageable even for a relatively small number of egress ports. Hence a separate multicast scheduler block is used in each ingress forwarder containing one real time multicast queue and one or more non-real time multicast queues as shown in Figure 1.

This structure of the ingress forwarder is highlighted in Figure 1. Incoming IP traffic from the line is queued in the relevant queues associated with a scheduler block. The choice of scheduler block is governed by the destination egress forwarder and whether it is multicast or unicast traffic. The class of service determines the specific queue to be utilised.

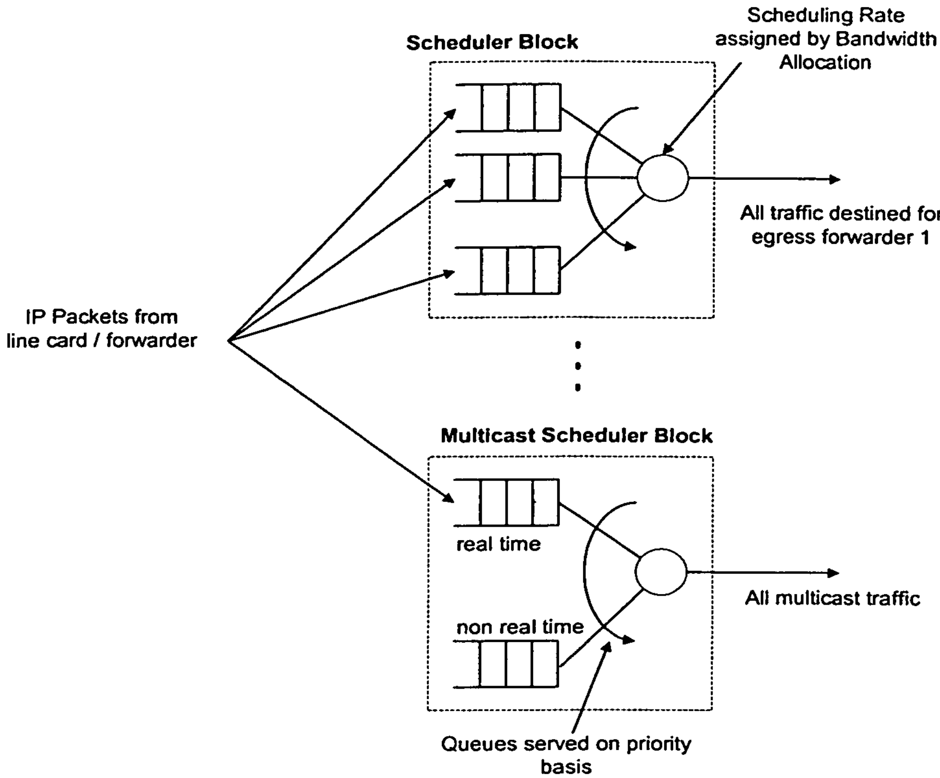

Figure 1 **Ingress Forwarder Scheduling Function**

The overall view of a centralised bandwidth allocation scheme is shown in Figure 2 below. Every fixed period (the Bandwidth Allocation Period – BAP) each ingress forwarder sends buffer occupancy (and possibly other information) to the central bandwidth allocation controller. In addition we also require each egress forwarder to send information on how many multicast cells were received in the last BAP from each ingress forwarder. The bandwidth allocation controller works out the allocation of bandwidth between all ingress/egress forwarder pairs for the next BAP and uses this to provide scheduling information to the ingress forwarders telling them which cells/packets to transmit in the next cell period.

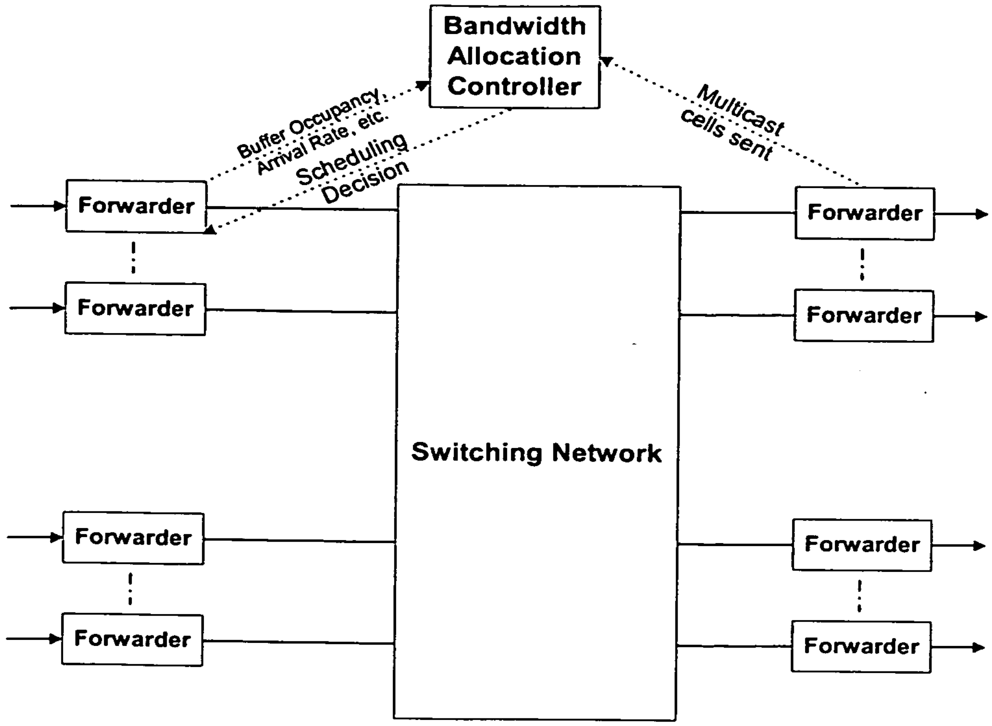

Figure 2 **Centralised Bandwidth Allocation**

In order to include multicast functionality into the bandwidth allocation process some additions are required to the unicast algorithm defined in the Roke Manor patent 1998P04863 (Centralised Bandwidth Allocation Mechanism). The unicast bandwidth allocation algorithm essentially divides the available bandwidth at ingress and egress amongst competing forwarders using the ingress queue length as a weighting factor. The queue length of the unicast scheduler block for traffic on ingress forwarder **i** destined for egress forwarder **j** is denoted by q_{ij}. Thus for example the amount of bandwidth allocated to unicast traffic from ingress forwarder **i** to egress forwarder **j** (be_{ij}) is given by the following equation.

$$be_{ij} = ABW_j * \frac{q_{ij}}{\sum_k q_{ik}} \qquad (1)$$

Here ABW_j is the available bandwidth (after real time reservations have been accounted for) at the egress forwarder j.

For real time multicast flows, the fan-out and bandwidth guarantees are known in advance and the sum of all ingress and egress requirements can be subtracted from the available bandwidth in the same way as for real time unicast traffic flows.

As the amount of egress bandwidth required for non-real time multicast flows is not known (compared with the case for real time multicast) it must be learnt by the system. One solution to this problem is to collect statistics at the egress forwarders on the number of multicast cells received from each ingress forwarder in the last Bandwidth Allocation Period (BAP). These statistics can then be included in the queue statistics message sent from the ingress forwarder to the central scheduler every BAP.

Figure 3 shows the relevant terminology used for calculating non-real time multicast bandwidth allocation. The ingress forwarder multicast queue occupancy is denoted as mcq_i for ingress forwarder **i**. The number of multicast cells received by egress forwarder **j** from ingress forwarder **i** in the last BAP is denoted by mcq_{ij}. The bandwidth allocated to non-real time multicast flows from ingress forwarder **i** to egress forwarder **j** is denoted by mcb_{ij}.

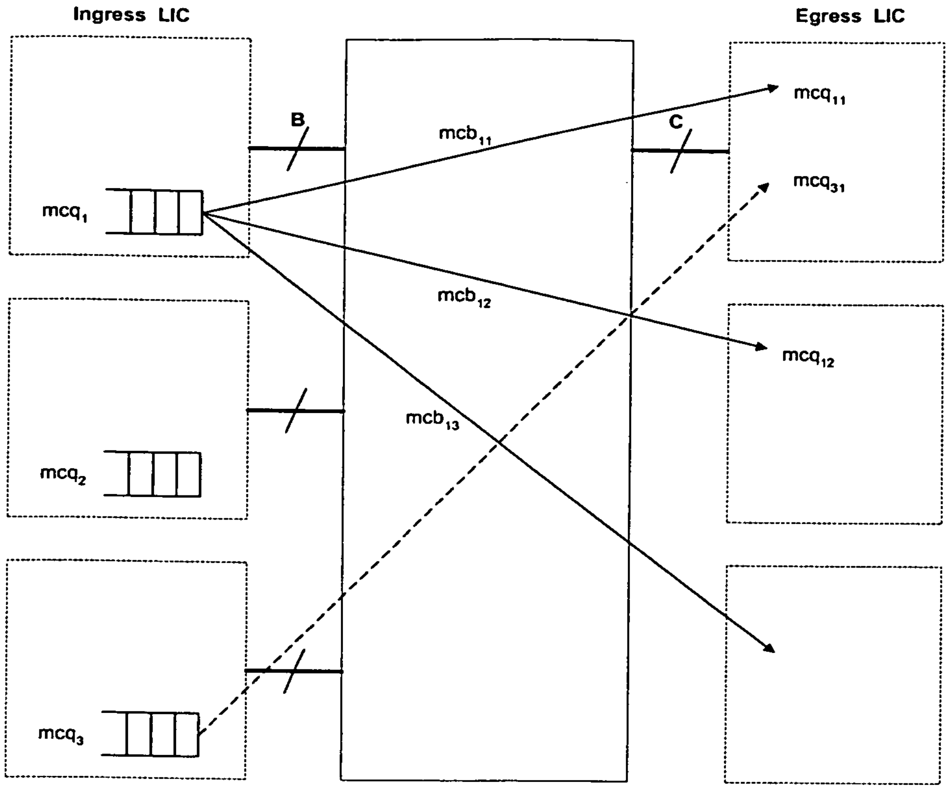

Figure 3 **Effective Queue lengths for Non-Real Time Multicast Bandwidth Allocation**

The value of mcq_i is used in the ingress bandwidth fair share in the same manner as q_{ij} does in the unicast centralised bandwidth allocation algorithm.

The values mcq_{ij} take part in the egress fair share allocation by providing a proportion with which to scale the ingress multicast queue occupancies. This means that the effective weight that the occupancy of the ingress nrt multicast queue (mcq_i) has on an egress forwarder j (called $emcq_{ij}$) is determined by the proportion of nrt cells received by egress forwarder j compared to all those sent in the last BAP period. It is therefore governed by the following equation:

$$emcq_{ij} = mcq_i * \frac{mcq_{ij}}{\sum_k mcq_{ik}} \qquad (2)$$

The value of $emcq_{ij}$ will be used in egress bandwidth allocation functions alongside the unicast equivalents q_{ij}.

Thus the equivalent of equation (1) when including the multicast traffic is given in equation (3).

$$bme_{ij} = ABW_j * \frac{emcq_{ij}}{\sum_k q_{kj} + \sum_k emcq_{kj}} \qquad (3)$$

Similar principles can be applied at the ingress for bandwidth allocation and any left over bandwidth can be distributed between unicast and multicast allocations by any fair mechanism required.